

HeteroQA: Learning towards Question-and-Answering through Multiple Information Sources via Heterogeneous Graph Modeling

Shen Gao¹, Yuchi Zhang², Yongliang Wang², Yang Dong², Xiuying Chen⁵,
Dongyan Zhao¹ and Rui Yan^{3,4*}

¹ Wangxuan Institute of Computer Technology, Peking University

² Ant Group

³ Gaoling School of Artificial Intelligence, Renmin University of China

⁴ Beijing Academy of Artificial Intelligence

⁵ King Abdullah University of Science and Technology

{shengao,zhaody}@pku.edu.cn, {yuchi.zyc,yongliang.wyl,doris.dy}@alibaba-inc.com

xiuying.chen@kaust.edu.sa,ruiyan@ruc.edu.cn

ABSTRACT

Community Question Answering (CQA) is a well-defined task that can be used in many scenarios, such as E-Commerce and online user community for special interests. In these communities, users can post articles, give comment, raise a question and answer it. These data form the heterogeneous information sources where each information source have their own special structure and context (comments attached to an article or related question with answers). Most of the CQA methods only incorporate articles or Wikipedia to extract knowledge and answer the user's question. However, various types of information sources in the community are not fully explored by these CQA methods and these multiple information sources (MIS) can provide more related knowledge to user's questions. Thus, we propose a question-aware heterogeneous graph transformer to incorporate the MIS in the user community to automatically generate the answer. To evaluate our proposed method, we conduct the experiments on two datasets: MSMP^{plus} the modified version of benchmark dataset MS-MARCO and the AntQA dataset which is the first large-scale CQA dataset with four types of MIS¹. Extensive experiments on two datasets show that our model outperforms all the baselines in terms of all the metrics.

CCS CONCEPTS

• Information systems → Question answering

KEYWORDS

Question answering system; Heterogeneous graph; Text generation

*Corresponding Author: Rui Yan (ruiyan@ruc.edu.cn) and Dongyan Zhao (zhaody@pku.edu.cn)

¹<https://github.com/gsh199449/HeteroQA>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '22, February 21–25, 2022, Tempe, AZ, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9132-0/22/02...\$15.00

<https://doi.org/10.1145/3488560.3498378>

ACM Reference Format:

Shen Gao, Yuchi Zhang, Yongliang Wang, Yang Dong, Xiuying Chen, Dongyan Zhao, & Rui Yan. 2022. HeteroQA: Learning towards Questionand-Answering through Multiple Information Sources via Heterogeneous Graph Modeling. In *Proceedings of the Fifteenth ACM Int'l Conference on Web Search and Data Mining (WSDM '22)*, Feb. 21–25, 2022, Tempe, AZ, USA. ACM, NY, NY, USA, 9 pages. <https://doi.org/10.1145/3488560.3498378>

1 INTRODUCTION

Online user community is widely adopted in many scenarios, such as software user communities (pytorch forum), interest groups (sub-reddit in Reddit and XueQiu), general communities (Quora and ZhiHu) and E-commerce websites (Taobao and AntFortune). To respond to user's queries timely and improve the experience, researchers propose the Community Question Answering (CQA) [11, 37, 47] which first reads the community article or Wikipedia and answers the user's question. Most previous works focus on multi-hop retrieving [33] or modeling the fine-grained interaction between question and documents [44].

However, most of the previous CQA models use a single information source (e.g., articles or knowledge base) and some of the information comes from an external source (e.g., Wikipedia), and it cannot cover the various aspects of user questions. In fact, there is various information in the online user community which has not been fully explored by existing CQA methods, such as article comments and related question-answer pairs. Moreover, the data structure and context of these multiple information sources (MIS) are very different: comments are associated with plain text articles and the related question-answering are pair-wise text data. Thus, these multiple information sources form a heterogeneous knowledge source which can help the CQA model improve the answer quality, as shown in Figure 1. The reason lies in that these information sources are supplements to the community articles and these sources from the same community are more relevant to the question than the retrieved external Wikipedia. Therefore, in this paper, we propose to incorporate four different types of MIS into the community question answering: (1) user articles; (2) article comments; (3) related questions; (4) their answers.

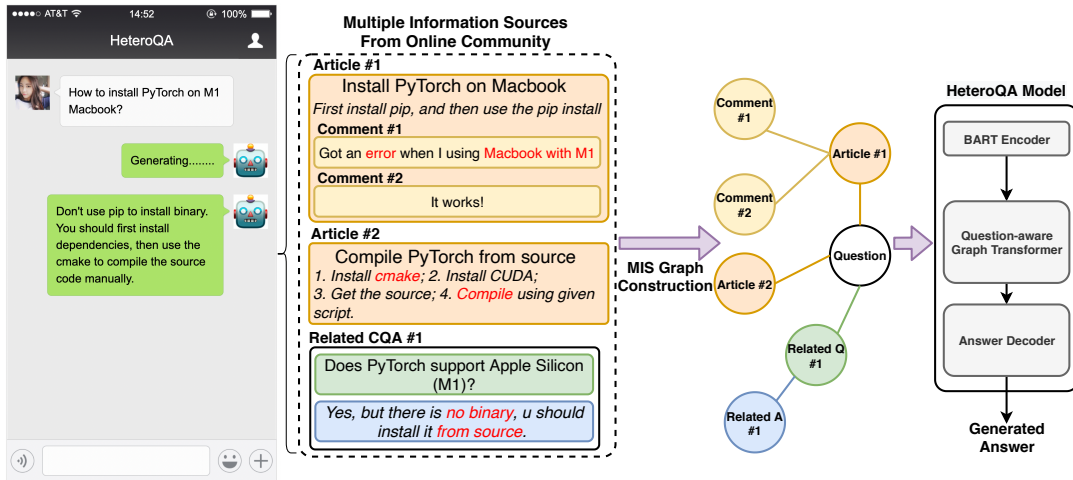


Figure 1: Example of incorporating heterogeneous information sources in community question answering.

Since MIS contains heterogeneous information form, we cannot simply treat them as plain text or model each type of MIS equally and independently, and each MIS should be modeled with its own structure and context. For example, the article comment cannot be understood without the semantic context of the article, and we should build an explicit connection between the comment and the article. To incorporate MIS into answer generation, there are two challenges that should be tackled: (1) How to unify the heterogeneous multiple information sources and extract useful knowledge according to the user question; (2) How to incorporate the knowledge learned from the heterogeneous MIS into the final answer generation process. Thus, we propose the **Heterogeneous graph based Question Answering model (HeteroQA)**. We first construct a heterogeneous graph using the different types of sources as a different type of node, and the graph construction method is shown in Figure 1. Then, we propose a question-aware graph reasoning method on the MIS heterogeneous graph to extract question-related knowledge from MIS. Finally, we generate the answer by dynamically incorporating graph node based on a BART-based [25] pretrained text generation framework. To demonstrate the effectiveness of the proposed HeteroQA model, we conduct experiments on two datasets: (1) MSM^{plus} dataset which is a modified version of the public benchmark CQA dataset MS-MARCO [3] and (2) a large-scale CQA dataset AntQA with four types of MIS that we released in this paper. Experiments demonstrate that HeteroQA outperforms all the strong baselines on both datasets.

To sum up, our contributions can be summarized as follows:

- We propose to incorporate MIS in the online user community to facilitate the CQA model.
- To leverage this heterogeneous information source, we propose the Heterogeneous graph-based Question Answering model (HeteroQA).
- Experiments show that our HeteroQA outperforms all the strong baselines on two datasets: MS-MARCO and AntQA which is a large-scale CQA dataset with MIS proposed in this paper.

2 RELATED WORK

Community Question Answering. The current CQA task aims to automate the process of finding appropriate answers to questions in a community-created discussion forum [34]. Recent researches [50, 53] often regard the CQA task as a text-matching problem and have proposed various deep learning networks to learn the semantic representation of question-answer pairs.

Different from the answer selection and span extraction methods, researchers propose the generative-based CQA methods [4] which can generate the natural language answers. Answering the custom question on E-commerce website is an important CQA application which has been addressed by many researchers [5, 15, 16], and most of these methods tackle this task as a knowledge-grounded text generation that leverages the custom reviews and product attributes to generate the answer.

However, existing CQA methods do not consider the MIS content in the community, which can provide more related knowledge than using the articles along to answer the question.

Heterogeneous Graph Network. Heterogeneous graph is widely used in many structure data modeling tasks: recommendation [8, 22, 30], social network modeling [9, 48, 51], text summarization [45], and knowledge-graph based tasks [21, 49]. In the CQA task, most of the researchers employ the homogeneous graph to model the contextual information, and only a few CQA methods leverage the heterogeneous graph model. Hu et al. [18, 19] propose to use the graph to model the multi-modal content of questions and answers with a graph neural network. Sun et al. [40] propose a graph-based model to route the newly posted questions to potential answers, which models the relationship between question, tags, and users using a heterogeneous graph. However, how to model the different types of MIS text using heterogeneous graph has not been studied in previous research, and most of the previous methods only retrieve existing answers or select answer from pre-defined candidates. In this paper, we focus on incorporating MIS from the user community and generate a proper answer from scratch.

Question: 在支付宝上借款怎么还款(How to conduct repayment on Alipay?)

Answer: 可以在支付宝里面进行还款(Use the Alipay mobile app to conduct the repayment.)

MIS (Related QA): Q: 支付宝借呗怎么还? (How to repay for "Jiebei" (Ant Microloan)?) **A:** 进入支付宝借呗页面, 就可以看到去借钱和去还钱, 点击还款就可以看到还款数目, 将足够的资金提前存入支付宝的余额(Login the "Jiebei" page, and you can see "Borrow money" and "Repayment". Click on the "Repayment" button, and then you can see the amount you should pay back. Deposit enough money in your Alipay balance account.)

MIS (Related Article): 借呗提前还款要收手续费吗? 不需要, 借呗提前还款暂时没有手续费, 大家直接在借呗页面操作即可。登录手机支付宝, 点击右下角, 在借呗首页选择, 按页面提示操作。(Do I have to charge a handling fee for early repayment? No, there is no handling fee for repayment in "Jiebei". So you can directly conduct repayment on the "Jiebei". First, log in to Alipay mobile version, click on the lower right corner, select on the home page of "Jiebei", and follow the instructions on the page.)

MIS (Related Article): 支付宝信用卡还款从3月26日起超出2000元部分将收取0.1%手续费的消息, 自昨天起在各大平台传得沸沸扬扬。相信大家其实都已经早有了心理准备。(The news that Alipay credit card repayments will be charged a 0.1% handling fee for more than 2,000 yuan from March 26 has been circulated on major platforms since yesterday. I believe that everyone has already been mentally prepared.) **Comments:** #1. 谢谢分享! (Thanks for sharing!) #2. 自从微信收费以来, 我一直是通过支付宝来还款(Since WeChat charges the handling fee, I have always used Alipay to repay)

Figure 2: Examples of the community QA and MIS.

Text Generation Pretraining. Recently, large-scale language models based on transformer [42] pretrained with mask language modeling [7] or text infilling [25] have been explored further advanced the state-of-the-art on many language understanding and generation tasks [27, 32]. We focus on improving the factual correctness and language fluency of the CQA model by using these pretraining models.

3 ANTQA DATASET

In this paper, we propose a new large-scale community question-and-answering (CQA) dataset named AntQA, and the biggest difference compared to existing CQA datasets lies in that AntQA contains multiple types of information sources to support the question-answering task. We will first introduce the data collection methods of AntQA, and then give some statistics of the QA pair and the multiple information sources.

3.1 Data Collection

To construct a community question-answering dataset, we collect the textual data from the AntFortune² which is one of the largest financial forums in China. In this financial forums, users in community discuss about the stock and mutual fund by publishing articles or sending tweets about their analysis of stock market. Other users can give comments to the articles and tweets. Most importantly, like most of online communities, users can pose a question and answer

²<https://antfortune.com>

others question. In this community, practitioners in the financial industry can obtain verification by showing their credentials to the administration of the online community. Thus, we collect the QA pairs which are answered by the verification users.

To build the multiple information sources (MIS), we collect four types of information sources including: (1) user articles; (2) article comments; (3) related questions and (4) its answers. For a community QA pair, we use the question as the query to retrieve the related question-answer from all the training QA samples using BM25 algorithm, and we use a popular information retrieval system ElasticSearch³ to implement distributed BM25 algorithm. And we also retrieve the user articles and its comments using the same method as the related QA pairs.

Table 2 shows a randomly sampled example data from the test set of AntQA. In this example, we show the question, answer and three MIS item which contains one related QA pair and two user articles. And the second article has two user comments. This example question asks about the usage of Alipay app, which topic is not a popular topic discussed by user articles. And how to repayment is a very common question especially for new users of Alipay, thus we can find the useful knowledge from the related QA pairs in MIS. However, not all the questions are talked about the simple topics, for example, other questions ask the trend of stock market should be answered by incorporating the user articles.

According to the user agreement and privacy policy of this online community, we are authorized to use and release the anonymous CQA and MIS data for research.

To verify the generalization of our proposed method in this paper, we also propose to modify a benchmark dataset MS-MARCO [3] as a new dataset MSM^{plus}. We employ the information retrieval system ElasticSearch to retrieve several similar QA pairs for each CQA sample and forming the new CQA dataset MSM^{plus} with related QA pairs as MIS. Since the MS-MARCO leaderboard is closed, we use the development set as the test set for MSM^{plus}. We will also introduce the statistics of MSM^{plus} in the next section for comparison with AntQA.

3.2 Key Statistics

We list some basic statistics for AntQA and MSM^{plus} in Table 1. AntQA and MSM^{plus} contain comparable number of training samples, and AntQA contains 376948 training samples, 1920 validation samples, and 2409 test samples. The AntQA contains more than three times supported articles than MSM^{plus} dataset, and AntQA has more related QA pairs than MSM^{plus}. In this paper, we use the character as the basic unit of Chinese and use the word unit for English. And the lengths of QA and articles in MSM^{plus} dataset are also shorter than the AntQA dataset.

To verify the correctness of the retrieved MIS in AntQA, we also conduct the human evaluation using scores 1-3 (3 means the MIS has completely correct information to answer the question). The average score is 2.12 and there are 72% samples that obtain a score equal to or greater than 2 which demonstrates the majority of MIS can support the model to give a correct answer.

³<https://www.elastic.co/cn/elasticsearch/>

Table 1: Dataset Statistics. The MIS are associated with each QA sample. The number of articles denotes the retrieved articles for a QA sample. The number of comments denotes the total number for the associated articles of a QA sample.

	AntQA	MSM ^{plus}
# of training samples	376948	387774
# of test samples	2409	49230
# of validation samples	1920	-
Avg. words of question	25.88	5.77
Avg. words of answer	83.89	16.27
Avg. number of articles	29.99	9.98
Avg. sentences of articles	3.87	5.04
Avg. words of articles	297.01	60.49
Avg. number of comments	11.63	-
Avg. words of comments	18.83	-
Avg. number of related QA	9.01	6.37

4 PROBLEM FORMULATION

Given a user question $q = \{w_1^q, \dots, w_{L_q}^q\}$ with L_q tokens and T types of related MIS $u = \{u^1, \dots, u^T\}$ which is retrieved from the database using the question as the query, our goal is to generate an answer $a = \{w_1^a, \dots, w_{L_a}^a\}$ with L_a tokens according to the knowledge learned from the related MIS. In each type of MIS u^i , it contains L_u^i documents $u^i = \{d_1^i, \dots, d_{L_u^i}^i\}$, and the document $d_j^i = \{d_{j,1}^i, \dots, d_{j,L_d}^i\}$ contains L_d tokens. In this paper, we retrieve four types of related MIS: (1) user articles; (2) article comments; (3) related questions and (4) its answers. Finally, we use the difference between generated answer \hat{a} and the ground truth answer a as the training objective.

5 HETEROQA MODEL

5.1 Overview

In this section, we introduce the **Heterogeneous graph based Question Answering model (HeteroQA)**. An overview of HeteroQA is shown in Figure 3, which has three main parts:

- **Node Encoder** employs the BART [25] to encode the question and each node text, and then uses the output hidden state of BART encoder as the initial graph node representation.
- **Question-aware Graph Transformer** encodes each type of MIS node and does reasoning among these nodes according to the question.
- **Answer Generator** generates the answer by incorporating question and updated graph nodes.

5.2 MIS Graph Construction

We first use the information retrieval system to retrieve related MIS as the knowledge source and construct a heterogeneous graph using these different types of MIS. We take each type of MIS as a different type of node, and the question as another type of node. Thus, we have five different types of nodes: (1) question; (2) user articles; (3) article comments; and (4) related questions and (5) its answers. Intuitively, we connect all types of MIS documents with

four basic edges: (1) user article to question; (2) related question to question; (3) related answer to related question; (4) comment to user article.

5.3 Node Encoder

Next, we employ the pretrained BART model to encode each MIS document independently:

$$\{\mathbf{h}_{j,1}^i, \dots, \mathbf{h}_{j,L_d}^i\} = \text{Enc}(d_{j,1}^i, \dots, d_{j,L_d}^i), \quad (1)$$

$$\{\mathbf{h}_1^q, \dots, \mathbf{h}_{L_q}^q\} = \text{Enc}(w_1^q, \dots, w_{L_q}^q), \quad (2)$$

where Enc is the encoder module in BART which outputs the vector representation $\mathbf{h}_{j,k}^i$ of each input token $d_{j,k}^i$ in j -th document of MIS type i . To obtain a vector representation of each MIS document, we apply the mean-pooling operation to the hidden states of tokens in a document:

$$\mathbf{h}_j^i = \text{Mean-pooling}(\{\mathbf{h}_{j,1}^i, \dots, \mathbf{h}_{j,L_d}^i\}). \quad (3)$$

We use the document representation \mathbf{h}_j^i as the initial node representation of the heterogeneous graph.

5.4 Question-aware Graph Transformer

The GNN (Graph Neural Network) has been proven to offer better performance for the reasoning task in QA [10, 21, 52]. The graph structure allows messages to pass over the nodes with local contextual information, which can comprehensively consider clue words in a different document to achieve the purpose of reasoning. Previous works in QA show that GNNs have achieved promising results in reasoning tasks [24, 31]. Inspired by the heterogeneous graph transformer [20], in this paper, we propose a Question-aware Graph Transformer (QGT) which can aggregate question-related information from different type of MIS node by passing information among heterogeneous nodes. An overview of the question-aware graph transformer is shown in the left part of Figure 3.

5.4.1 Heterogeneous Question-aware Attention. Since the structure information stored in the MIS graph is important for capturing the relationship between different types of MIS, we should fuse the information in neighbor nodes into the target node representation. However, the salience of each node is not the same, and it is unreasonable to simply add all the information of all neighbor nodes into the target node. Therefore, we conduct the multi-head mechanism [42] to calculate the relevance score between the target node and the neighbor nodes. We first project the input node representation into two spaces using different parameters for different type of MIS node:

$$V(s) = \text{MLP}_{\gamma(s)}(D[s]), \quad (4)$$

$$K(t) = \text{MLP}_{\gamma(t)}(D[t]), \quad (5)$$

where $\gamma(\cdot)$ denotes the type of node, $D[\cdot]$ denotes the original node representation, s denotes the index of one source node and t denotes the index of the target node. Next, we employ a bi-linear function to calculate the relevance between these nodes as the attention scores, and we also incorporate the edge type into the attention score calculation:

$$\alpha(s, e, t) = \text{Softmax}_{\forall s \in N(t)}(K(t)W_{\delta(e)}V(s)), \quad (6)$$

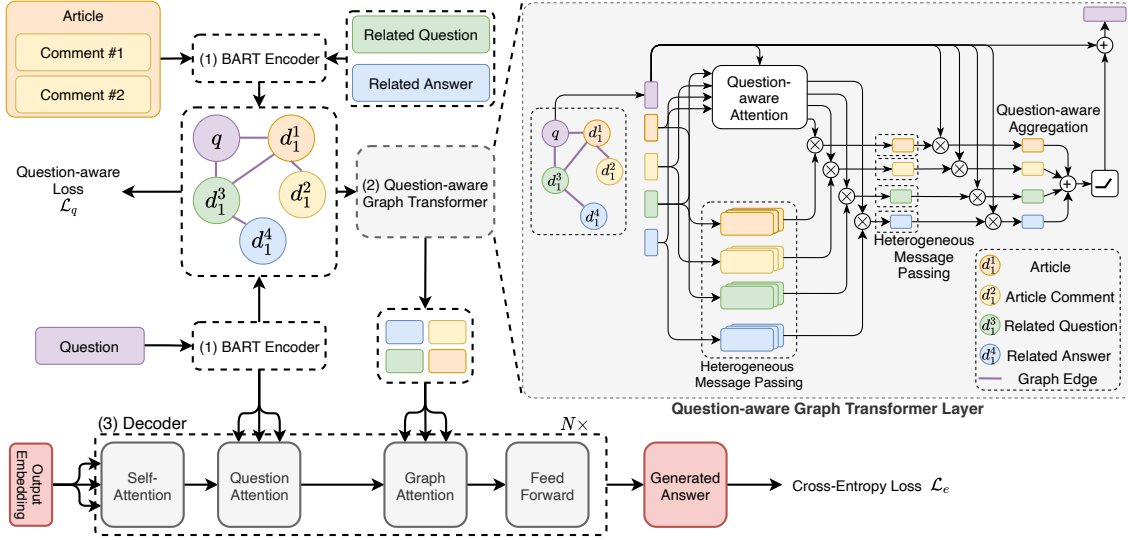


Figure 3: Overview of HeteroQA. We divide our model into three parts: (1) BART-based Node Encoder; (2) Question-aware Graph Transformer; (3) Answer Decoder.

where e is the edge between the node s and node t , $N(t)$ is the neighbour nodes of node t , and the $W_{\delta(e)}$ is the edge type specific parameter.

5.4.2 Heterogeneous Message Passing. Intuitively, the message passing behavior on different edge types should not be the same. For example, when passing a message between two user articles connected by an edge with the same entity words, we should pass the facts about the entity. And the message between the related question and its answer should be the extracted facts from the answer according to the related answer. Thus, the passed message from one node to other target nodes should not be the same when passing through a different type of edge. We need to extract the useful message from the source node by considering the edge type:

$$M(s, e, t) = \text{MLP}_{\gamma(s)}^{\text{msg}}(D[s]) W_{\delta(e)}^{\text{msg}}, \quad (7)$$

where $W_{\delta(e)}^{\text{msg}}$ is the edge type specific parameter matrix, and $M(s, e, t)$ is the message of source node s that will be passed through edge e to the target node t .

5.4.3 Question-aware Aggregation. Since the aim of conducting the message passing between nodes is to extract useful knowledge from heterogeneous MIS to answer the question, only the question-related knowledge should be reserved when passing the message. Thus, first use the question relation score to re-scale the attention scores (calculated in Equation 6):

$$\hat{\alpha}(s, e, t) = \alpha(s, e, t) \times \beta(s), \quad (8)$$

where $\beta(s) \in \mathbb{R}$ is the score for node s which represents the semantic relevance between node s and question q calculated by a bi-linear layer:

$$\beta(s) = D[q] W^r D[s], \quad (9)$$

where W^r is a trainable parameter, and the $D[q]$ and $D[s]$ are node representations for question q and source node s respectively. Then,

we apply a weighted sum over all the messages passed from all the source nodes:

$$\tilde{D}[t] = \sum_{\forall s \in N(t)} (\hat{\alpha}(s, e, t) \cdot M(s, e, t)). \quad (10)$$

Finally, we apply a linear projection with non-linear activation and a residual connect on the source nodes messages:

$$\hat{D}[t] = \text{gelu}(\text{MLP}_{\gamma(t)} \tilde{D}[t]) + D[t], \quad (11)$$

where $\hat{D}[t]$ is the updated node representation for target node t . In this way, we obtain the contextualized vector representation of the target node which fuses the question-related information from neighborhood MIS nodes.

Inspired by the multi-layer mechanism of Transformer [42], we also conduct this question-aware graph transformer several times. In each layer, the input is the output of the previous layer and the output will be used in the next layer as input. The first layer uses the initial node representation as the input which is illustrated in § 5.3. After using the multi-layer heterogeneous graph transformer to encode the target node in the graph, we obtain the final vector representation of the target node. For brevity, we omit the layer notion in our notion and equations, and use $\mathbf{g}_j^t = \hat{D}[t]$ as the notion of final node representation for the node t which is the j -th node in i -th MIS type.

5.5 Answer Generator

HeteroQA generates the answer based on the knowledge learned from the MIS. We employ the pretrained BART decoder as our base framework. However, generating a good answer should incorporate the input question and the knowledge from MIS simultaneously, and the original BART decoder can only attend to the hidden states produced by the BART encoder. We introduce a new graph attention layer and insert this layer into the original BART decoder, and

preserve other layers in the BART. We first apply the self-attention on the masked output embeddings of previous decoding steps which is same as the Transformer [42], and we obtain the output \mathbf{p}^s for each decoding step. Then we use the output \mathbf{p}^s to cross-attend to the question hidden states:

$$\mathbf{p}^q = \text{MHAtt}(\mathbf{p}^s, \{\mathbf{h}_1^q, \dots, \mathbf{h}_{L_q}^q\}), \quad (12)$$

where MHAtt is the standard multi-head attention layer [42] and this procedure is the same as the original BART decoder. After the self-attention and cross-attention layer, we apply a multi-head attention layer which aggregate useful knowledge from the updated graph nodes according to the state of each decoding step:

$$\mathbf{p}^g = \text{MHAtt}(\mathbf{p}^s, \{\mathbf{g}_1^1, \dots, \mathbf{g}_{L_T}^T\}). \quad (13)$$

Finally, we combine the question and knowledge from MIS:

$$\mathbf{p}^f = \mathbf{p}^q + \mathbf{p}^g. \quad (14)$$

We use the fused representation \mathbf{p}^f to predict the distribution over the vocabulary of the generated answer, and use the cross-entropy loss \mathcal{L}_e as the training objective to optimize the parameters.

5.5.1 Loss Function. To facilitate the question-aware graph transformer to capture the semantic relationship between each retrieved MIS and question, we propose an auxiliary question-aware loss function. In detail, we employ a linear layer that uses the updated node representation as input to predict the BM25 score of each MIS node, and we calculate the mean squared error between the predicted score and the BM25 score produced by the information retrieval system as an auxiliary loss \mathcal{L}_q . Finally, we combine text generation cross-entropy loss \mathcal{L}_e and question-aware loss \mathcal{L}_q as the final loss function $\mathcal{L} = \mathcal{L}_e + \psi \mathcal{L}_q$ where ψ is a hyper-parameter.

6 EXPERIMENTAL SETUP

6.1 Research questions

We list four research questions that guide the remainder of the paper:

- (1) **RQ1** (See §7.1): What is the overall performance of HeteroQA?
- (2) **RQ2** (See §7.2): What is the effect of each module in HeteroQA?
- (3) **RQ3** (See §7.3): Does every type of MIS help the CQA model to generate better answer?
- (4) **RQ4** (See §7.4): What is the effect of using more or less MIS in each type?

6.2 Evaluation Metrics

To evaluate the methods, following existing text generation methods [12–14], we employ the automatic word overlap-based BLEU [35], METEOR [2] and ROUGE-L [29] to measure the lexical unit overlapping (e.g., unigram, bigram) between the generated answer and ground truth. Since only using automatic evaluation metrics can be misleading [39], we also conduct the human evaluation. Three well-educated and full-time professional annotators are invited to judge 200 randomly sampled answers. The statistical significance of two runs is tested using a two-tailed paired t-test and is denoted using \blacktriangle (or \blacktriangledown) for strong significance for $\alpha = 0.01$.

6.3 Comparisons

To prove the effectiveness of each module, we conduct ablation studies which remove each key module in HeteroQA, and then form three baseline methods (1) HeteroQA-QGT uses the original HGT [20] instead of our proposed question-graph transformer; (2) HeteroQA-GL removes the graph auxiliary loss \mathcal{L}_q ; (3) HeteroQA-HG uses the homogeneous graph model GAT [43] as the graph encoder. Apart from the ablation study, we also compare with the following baselines:

- Retrieved1 uses the top-1 retrieved document as the answer for the user question.
- S2SA is the Sequence-to-Sequence framework [41] which is equipped with the attention mechanism [1] and copy mechanism [17, 38] as a baseline method. The input sequence is a question and the ground truth output sequence is the answer.
- S2SAU is a simple method that can incorporate the MIS information when generating the answer, we concatenate all types of MIS after the question as the input of S2SA.
- CVAE [54] is a text generation framework based on conditional variational autoencoder that uses latent variables to generate diverse sentences, and this framework has been applied in many knowledge-grounded text generation tasks [23, 28].
- BART [25] is a denoising autoencoder for pretraining sequence-to-sequence models, and it is particularly effective when fine-tuned for text generation but also works well for comprehension tasks.
- RAG [26] is general-purpose fine-tuning recipe for retrieval-augmented generation models which combine pretrained non-parametric memory for knowledge-intensive NLP tasks. RAG outperforms many strong Open-Domain QA and CQA methods.
- OAG [6] is a generative community question answering model which jointly models opinionated and interrelated information between the question and contextual document.

6.4 Implementation Details

We implement our experiments using PyTorch [36] based on the Transformers [46]. We train our model on two NVIDIA V100 GPU for four days. We employ the pretrained BART-base model (with 6 layers for encoder and decoder, the number of attention head is 12, and the hidden size is 768) to initialize part of the parameters. The hyper-parameter ψ is set to 0.01. Since there is not a public pretrained Chinese BART, we pretrain the BART-base from scratch using the text infilling objective on 7GB Chinese text data, and train the model for two weeks using 8 NVIDIA V100 GPU.

7 EXPERIMENTAL RESULTS

7.1 Overall Performance

We compare our model with the baselines listed in Table 2. Our model performs consistently better on two datasets than other CQA and conditional text generation models with improvements of 13.73%, 9.68%, and 13.22% on the AntQA dataset, and achieves 33.09%, 12.00% and 27.63% improvements on the MSMP^{plus} dataset compared with RAG in terms of BLEU, BLEU1, and BLEU2 respectively. This demonstrates that MIS in the user community provides useful knowledge for answer generation that cannot be replaced by other complicated baselines without this information. And our

Table 2: Automatic metrics comparison between baselines.

Method	AntQA							MSM ^{plus}						
	BLEU	BLEU1	BLEU2	BLEU3	BLEU4	ROUGE-L	METEOR	BLEU	BLEU1	BLEU2	BLEU3	BLEU4	ROUGE-L	METEOR
Retrieved1	1.11	8.35	1.95	0.52	0.18	7.32	0.15	0.21	4.82	0.25	0.04	0.03	3.78	0.07
S2SA	13.40	24.54	14.01	10.57	8.88	19.42	0.18	0.39	7.36	0.82	0.13	0.03	10.08	0.06
SASAU	14.80	25.14	15.14	11.98	10.53	19.73	0.18	0.57	8.60	1.09	0.20	0.05	11.23	0.07
CVAE	13.15	23.10	13.54	10.55	9.07	18.44	0.17	0.38	7.45	0.80	0.13	0.03	10.37	0.05
OAAG	2.19	3.77	2.32	1.77	1.49	17.07	0.13	1.56	8.61	1.25	0.78	0.70	8.49	0.06
BART	16.47	27.93	16.99	13.37	11.61	22.06	0.21	6.66	21.06	8.28	4.34	2.60	28.82	0.19
RAG	17.33	29.04	17.93	14.11	12.27	22.51	0.21	6.95	21.58	8.76	4.60	2.68	29.47	0.23
HeteroQA	19.71	31.85	20.30	16.28	14.35	22.78	0.22	9.25	24.17	11.18	6.50	4.17	34.63	0.28
HeteroQA-QGT	18.31	29.97	18.83	15.04	13.25	21.84	0.20	7.18	21.10	8.85	4.83	2.95	32.90	0.26
HeteroQA-GL	18.90	30.80	19.46	15.53	13.71	22.26	0.20	7.20	21.06	8.85	4.85	2.97	32.96	0.26
HeteroQA-HG	18.20	30.29	18.79	14.87	12.98	21.78	0.20	5.36	17.93	6.69	3.42	2.01	28.40	0.22

Table 3: Human Evaluation Results.

Dataset	Method	Fluency	Correctness
AntQA	OAAG	2.49	1.84
	BART	2.58	1.96
	RAG	2.44	2.11
	HeteroQA	2.61[▲]	2.25[▲]
MSM ^{plus}	OAAG	2.68	1.56
	BART	2.89	1.65
	RAG	2.88	1.55
	HeteroQA	2.90[▲]	1.83[▲]

HeteroQA method has a greater performance improvement compared to the methods that do not use the pretrained language model. Sun et al. [40] also propose a retrieval-based CQA method, which leverages heterogeneous graph to route the cold question to the exiting QA nodes in the CQA graph and they did not incorporate the MIS in the user community. Since they did not release their source code, we use a simple and intuitive method Retrieved1 to verify the basic performance of these retrieval-based methods. From Table 2 we can find that the retrieval-based method achieves the worse performance among all the baselines, which demonstrates the retrieval-based method is not suitable for the generative CQA.

For the human evaluation, we asked annotators to rate each summary according to its factual correctness and language fluency. The rating score ranges from 1 to 3, with 3 being the best. Table 3 lists the average scores of each model, showing that HeteroQA outperforms the other baseline models in both fluency and correctness. To prove the significance of these results, we also conduct the paired student t-test between our model and RAG and obtain $p < 0.05$ for fluency and consistency, respectively. From this experiment, we find that the BART achieves a comparable fluency score as HeteroQA, the correctness score is significantly lower than HeteroQA. We also find that there are 22% samples generated by our model obtain a low score (1) for correctness. To verify the reason for generating these bad cases, we find that in the input MIS of these bad cases, there are 86.36% samples that obtain a score 1 for the quality of MIS. This phenomenon demonstrates that the misleading MIS input leads our HeteroQA model to generate an answer with wrong facts.

Table 4: Effective of using different type of MIS.

Dataset	Method	BLEU	ROUGE-L	METEOR
AntQA	HeteroQA	19.71	22.78	0.22
	HeteroQA-QA	17.76	20.95	0.20
	HeteroQA-COM	18.56	21.81	0.21
	HeteroQA-ART	17.91	21.55	0.19
MSM ^{plus}	HeteroQA	9.25	34.63	0.28
	HeteroQA-QA	7.08	32.73	0.25
	HeteroQA-COM	-	-	-
	HeteroQA-ART	5.31	28.47	0.22

Although our HeteroQA model achieves 13.73% improvement in terms of BLEU over RAG model on AntQA, the increment is smaller than the performance on the MSM^{plus} dataset which achieves 33.09% improvement in terms of BLEU. The reason is that AntQA is more similar to the real-world scenario where the supported MIS are all retrieved by an information retrieval system using the BM25 algorithm and we did not employ any labor-intensive human annotation to filter the data which contains some noisy samples. Conversely, the supporting document in MSM^{plus} is filtered by a human annotator, which has less noisy data in their dataset.

7.2 Ablation Studies

The performance of different ablation models are shown in Table 2. All ablation models perform worse than HeteroQA in terms of all metrics, which demonstrates the preeminence of HeteroQA. The experiment shows that our HeteroQA achieves 7.65% and 28.83% increments compared with HeteroQA-QGT in terms of BLEU score on AntQA and MSM^{plus} respectively, which verifies the effectiveness of our proposed QGT. In the method HeteroQA-HG, we employ the GAT [43] as the graph encoder and there is a slight decrement from HeteroQA-HG to HeteroQA-QGT of 0.28% and 15.85% in terms of ROUGE-L on AntQA and MSM^{plus}, which demonstrates the effectiveness of treating each type of MIS as different nodes and using the heterogeneous graph to model these MIS. Thus, we can verify the contributions of each module in HeteroQA.

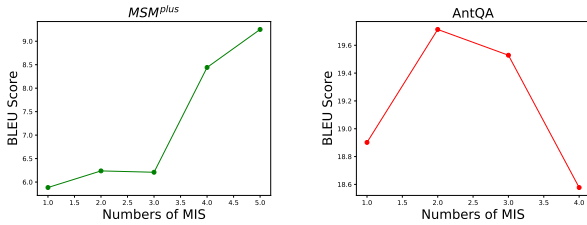


Figure 4: Visualization of BLEU score with different MIS document number.

7.3 Effect of Different Type of MIS

To verify the effectiveness of each MIS type, we remove each type of MIS from our heterogeneous graph. We employ three experiments that remove the MIS type one-by-one: HeteroQA-QA (w/ comment, article), HeteroQA-COM (w/ article, related question and answer) and HeteroQA-ART (w/ related question and answer). Since the comment of the documents in MSM^{plus} can not be obtained, we omit this experiment. Table 4 shows the performance of these ablation models, and we only show three metrics of each model due to the limited space. From this table, we find that the related QA produces the greatest contribution (10.98% decline in terms of the BLEU) on AntQA, and the related articles produce the greatest contribution on MSM^{plus} with a 74.20% decline in terms of the BLEU. Although the contribution of using comment of articles is minimum compared to other MIS types, it also increases the BLEU score for 6.20%.

7.4 Effect of Using Different Number of MIS

We investigate the influence of the document number for each MIS type incorporated by the HeteroQA model. Figure 4 illustrates the relationship between the document number in each type of MIS and the BLEU score. On the MSM^{plus} dataset, the performance keeps growing when the number of document reaches 5, and due to the hardware limitation, we did not conduct experiments for more MIS. Since the document in MSM^{plus} is filtered by humans and all the related articles can support answering the question, the performance of the model can keep increasing when the model reads more MIS. On the AntQA dataset, the BLEU score first increases with the number of MIS, and after reaching an upper limit it then begins to drop. This phenomenon demonstrates that incorporating several relevant MIS can help the model to achieve better performance. In the real-world application, we cannot filter the contextual information case-by-case. And excessive MIS which contains not relevant information will confuse the model to extract useful knowledge to generate correct answers.

7.5 Case Study

Table 5 shows an example of the generated answer and MIS context with useful knowledge to answer the user question. Although all the methods generate fluent answers, the facts described in BART and RAG are wrong (the wrong facts are shown in red text in Table 5). And HeteroQA generates refined and correct repayment steps for the user question according to the related QA and user article.

Question: 在支付宝上借款怎么还款(How to conduct repayment on Alipay?)

MIS (Related QA): Q: 支付宝借呗怎么还? (How to repay for "Jiebei" (Ant Microloan)?) A: 进入支付宝借呗页面, 就可以看到去借钱和去还钱, 点击还款就可以看到还款数目, 将足够的资金提前存入支付宝的余额(Login the "Jiebei" page, and you can see "Borrow money" and "Repayment". Click on the "Repayment" button, and then you can see the amount you should pay back. Deposit enough money in your Alipay balance account.)

MIS (Related Article): 借呗提前还款要收手续费吗? 不需要, 借呗提前还款暂时没有手续费, 大家直接在借呗页面操作即可。登录手机支付宝, 点击右下角, 在借呗首页选择, 按页面提示操作。(Do I have to charge a handling fee for early repayment? No, there is no handling fee for repayment in "Jiebei". So you can directly conduct repayment on the "Jiebei". First, log in to Alipay mobile version, click on the lower right corner, select on the home page of "Jiebei", and follow the instructions on the page.)

BART: 登录支付宝, 点击我的, 点击借款, 点击还款就可以了(Login Alipay, then click on "My Apps", then click on "Borrow some money", then click on "Repayment" and that's it.)

RAG: 支付宝借呗借款的步骤, 打开支付宝钱包, 点击借呗, 然后点击还款, 选择借款方式, 就可以还款了(The steps of borrowing money from "Jiebei" (Ant Microloan). Login "Pocket" and click on "Jiebei". Then click on "Repayment" and select the way of making a loan.)

HAG: 进入支付宝, 点击借呗, 然后点击还款就可以了(Login Alipay, click on "Jiebei" (Ant Microloan), and then click on "Repayment".)

Figure 5: Examples of the context and answers. We only show two MIS due to the limited space. We use red text to indicate wrong facts in generated answers.

8 CONCLUSION

In this paper, we propose the Heterogeneous graph-based Question Answering model (HeteroQA) to tackle the community question answering task, which extracts useful knowledge from Multiple Information Sources (MIS) which contains four types documents in the user community. Since MIS have different contexts and structures, we treat MIS as heterogeneous information and propose a question-aware heterogeneous graph transformer to model the relationship between these nodes. To generate correct and fluent answer text, we leverage the pretrain language model as an answer generator and introduce a graph attention layer to fuse the knowledge into the answer generation. Experiments conducted on two datasets verify the effectiveness of our HeteroQA model.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their constructive comments. We would also like to thank Nan Su for her help on this project. This work was supported by the National Key Research and Development Program of China (No. 2020AAA0106600), the National Science Foundation of China (NSFC No. 62122089 and No. 61876196). Rui Yan is supported as a young fellow at Beijing Academy of Artificial Intelligence (BAAI).

REFERENCES

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *CoRR*.
- [2] Satantjeve Banerjee and Alon Lavie. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*. 65–72.
- [3] Daniel Fernando Campos, T. Nguyen, M. Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, L. Deng, and Bhaskar Mitra. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. *ArXiv abs/1611.09268* (2016).
- [4] Yang Deng, Wai Lam, Yuxiang Xie, Daoyuan Chen, Yaliang Li, Min Yang, and Ying Shen. 2020. Joint Learning of Answer Selection and Answer Summary Generation in Community Question Answering. In *AAAI*.
- [5] Yang Deng, Wenxuan Zhang, and Wai Lam. 2020. Opinion-aware Answer Generation for Review-driven Question Answering in E-Commerce. In *Arxiv*.
- [6] Yang Deng, Wenxuan Zhang, and Wai Lam. 2020. Opinion-aware Answer Generation for Review-driven Question Answering in E-Commerce. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*. 255–264. <https://doi.org/10.1145/3340531.3411904>
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*.
- [8] Shaohua Fan, Junxiong Zhu, Xiaotian Han, Chuan Shi, Linmei Hu, Biyu Ma, and Yongliang Li. 2019. Metapath-guided Heterogeneous Graph Neural Network for Intent Recommendation. In *KDD*.
- [9] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph Neural Networks for Social Recommendation. In *WWW*.
- [10] Yuwei Fang, Siqi Sun, Zhe Gan, Rohit Pillai, Shuohang Wang, and Jingjing Liu. 2020. Hierarchical Graph Network for Multi-hop Question Answering. In *EMNLP*.
- [11] Jinlan Fu, Yi Li, Qi Zhang, Qinzhuo Wu, Renfeng Ma, Xuanjing Huang, and Yugang Jiang. 2020. Recurrent Memory Reasoning Network for Expert Finding in Community Question Answering. In *WSDM*.
- [12] Shen Gao, Xiuying Chen, Piji Li, Zhangming Chan, Dongyan Zhao, and Rui Yan. 2019. How to Write Summaries with Patterns? Learning towards Abstractive Summarization through Prototype Editing. In *EMNLP*.
- [13] Shen Gao, Xiuying Chen, Piji Li, Zhaochun Ren, Lidong Bing, Dongyan Zhao, and Rui Yan. 2019. Abstractive Text Summarization by Incorporating Reader Comments. In *AAAI*.
- [14] Shen Gao, Xiuying Chen, Zhaochun Ren, Dongyan Zhao, and Rui Yan. 2020. From Standard Summarization to New Tasks and Beyond: Summarization with Manifold Information. In *IJCAI*.
- [15] Shen Gao, Xiuying Chen, Zhaochun Ren, Dongyan Zhao, and Rui Yan. 2021. Meaningful Answer Generation of E-Commerce Question-Answering. *ACM Trans. Inf. Syst.* 39, 2, Article 18 (Feb. 2021), 26 pages.
- [16] Shen Gao, Zhaochun Ren, Yihong Zhao, Dongyan Zhao, Dawei Yin, and Rui Yan. 2019. Product-Aware Answer Generation in E-Commerce Question-Answering. In *WSDM*.
- [17] Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O. K. Li. 2016. Incorporating Copying Mechanism in Sequence-to-Sequence Learning. *CoRR abs/1603.06393* (2016).
- [18] Jun Hu, Quan Fang, Shengsheng Qian, and Changsheng Xu. 2020. Multi-modal Attentive Graph Pooling Model for Community Question Answer Matching. In *MM*.
- [19] Jun Hu, Shengsheng Qian, Quan Fang, and Changsheng Xu. 2019. Hierarchical Graph Semantic Pooling Network for Multi-modal Community Question Answer Matching. In *MM*.
- [20] Ziniu Hu, Yuxiao Dong, Kuansan Wang, Yizhou Sun, Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous Graph Transformer. In *WWW*.
- [21] Pin Jiang and Yahong Han. 2020. Reasoning with Heterogeneous Graph Alignment for Video Question Answering. In *AAAI*.
- [22] Zhuoren Jiang, Yue Yin, Liangcai Gao, Yao Lu, and Xiaozhong Liu. 2018. Cross-language Citation Recommendation via Hierarchical Representation Learning on Heterogeneous Graph. In *SIGIR*.
- [23] Byeongchang Kim, Jaewoo Ahn, and Gunhee Kim. 2020. Sequential Latent Knowledge Selection for Knowledge-Grounded Dialogue. In *ICLR*.
- [24] Yunshi Lan and Jing Jiang. 2020. Query Graph Generation for Answering Multi-hop Complex Questions from Knowledge Bases. In *ACL*.
- [25] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL*.
- [26] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen tau Yih, Tim Rocktaschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *NIPS*.
- [27] Jinpeng Li, Yingce Xia, Rui Yan, Hongda Sun, Dongyan Zhao, and Tie-Yan Liu. 2021. Stylized Dialogue Generation with Multi-Pass Dual Learning. In *NeurIPS*.
- [28] Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to Select Knowledge for Response Generation in Dialog Systems. In *IJCAI*.
- [29] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. *Text Summarization Branches Out* (2004).
- [30] Siwei Liu, Iadh Ounis, Craig Macdonald, and Zaiqiao Meng. 2020. A Heterogeneous Graph Neural Model for Cold-start Recommendation. In *SIGIR*.
- [31] Shangwen Lv, Daya Guo, Jingjing Xu, Duyu Tang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, and Songlin Hu. 2020. Graph-Based Reasoning over Heterogeneous External Knowledge for Commonsense Question Answering. In *AAAI*.
- [32] Yougang Lyu, Zihan Wang, Zhaochun Ren, Pengjie Ren, Zhumin Chen, Xiaozhong Liu, Yujun Li, Hongsong Li, and Hongye Song. 2022. Improving legal judgment prediction through reinforced criminal element extraction. *Information Processing & Management* 59, 1 (2022), 102780.
- [33] Sewon Min, Danqi Chen, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2020. Knowledge Guided Text Retrieval and Reading for Open Domain Question Answering. In *Arxiv*.
- [34] Preslav Nakov, Lluís Márquez i Villodre, Alessandro Moschitti, Walid Magdy, Hamdy Mubarak, Abed Alhakim Freihat, James R. Glass, and Bilal Randeree. 2016. SemEval-2016 Task 3: Community Question Answering. *ArXiv abs/1912.01972* (2016).
- [35] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *ACL*.
- [36] Adam Paszke, S. Gross, and et. al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*.
- [37] Andreas Ruckle, Krishnkant Swarnkar, and Iryna Gurevych. 2019. Improved Cross-Lingual Question Retrieval for Community Question Answering. In *WWW*.
- [38] Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *ACL*.
- [39] Amanda Stent, Matthew Marge, and Mohit Singhai. 2005. Evaluating Evaluation Methods for Generation in the Presence of Variation. In *CICLING*.
- [40] Jiankai Sun, Jie Zhao, Huan Sun, and Srinivasan Parthasarathy. 2020. EndCold: An End-to-End Framework for Cold Question Routing in Community Question Answering Services. In *IJCAI*.
- [41] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to Sequence Learning with Neural Networks. In *NIPS*.
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*.
- [43] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR*.
- [44] Bingning Wang, Ting Yao, Qi Zhang, Jingfang Xu, Zhixing Tian, Kang Liu, and Jun Zhao. 2019. Document Gated Reader for Open-Domain Question Answering. In *SIGIR*.
- [45] Danqing Wang, Pengfei Liu, Yining Zheng, Xipeng Qiu, and Xuanjing Huang. 2020. Heterogeneous Graph Neural Networks for Extractive Document Summarization. In *ACL*.
- [46] Thomas Wolf, Lysandre Debut, and et. al. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45.
- [47] Wei Wu, Xu Sun, and Houfeng Wang. 2018. Question Condensing Networks for Answer Selection in Community Question Answering. In *ACL*.
- [48] Yongji Wu, Defu Lian, Shuwei Jin, and Enhong Chen. 2019. Graph Convolutional Networks on User Mobility Heterogeneous Graphs for Social Relationship Inference. In *IJCAI*.
- [49] Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, Rui Yan, and Dongyan Zhao. 2019. Relation-Aware Entity Alignment for Heterogeneous Knowledge Graphs. In *IJCAI*.
- [50] Yuxiang Xie, Ying Shen, Yaliang Li, Min Yang, and Kai Lei. 2020. Attentive User-Engaged Adversarial Neural Network for Community Question Answering. In *AAAI*.
- [51] Xiaoyu Yang, Yuefei Lyu, Tian Tian, Yifei Liu, Yudong Liu, and Xi Zhang. 2020. Rumor Detection on Social Media with Graph Structured Adversarial Learning. In *IJCAI*.
- [52] Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, and Jure Leskovec. 2021. QA-GNN: Reasoning with Language Models and Knowledge Graphs for Question Answering. In *NAACL*.
- [53] Xiaodong Zhang, Sujian Li, Lei Sha, and Houfeng Wang. 2017. Attentive Interactive Neural Networks for Answer Selection in Community Question Answering. In *AAAI*.
- [54] Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning Discourse-level Diversity for Neural Dialog Models using Conditional Variational Autoencoders. In *ACL*.